10/521729
DT1 Rec'd PCT/PTO 20 JAN 2005

3/PRTS

Image Control

This invention relates to control systems employing images produced by radiation-generating imaging plant or equipment. It is particularly concerned with medical and industrial uses of radiation imaging and especially with X-ray imaging.

A huge variety of imaging systems are known and well established. Their aim is to produce a tangible representation of a target under inspection. In addition to X-ray and optical images they include thermal (IR), magnetic resonance, radio frequency (radar) and ultrasound images and combinations of two or more of these. Their use however has been mostly limited to the generation of the image and to its study by suitably trained observers or operators who may employ the image data for purposes of diagnostics or process control.

It is known to employ sensors, either single sensors or collections of point sensors, for such process control duties as thickness gauging. There have also been proposals to use an image parameter to control an aspect of medical or industrial systems. In the field of medical X-ray imaging, US patent specification No. 5,253,169 proposes the use of a collimator which moves according to the location of a monitored catheter tip, but gives no details of the necessary hardware nor of any filtering function for controlling the system. US patent specification No. 5,278,887 proposes for medical X-ray imaging the use of semitransparent collimators that move automatically in response to a "medical instrument" but it has no details of the control strategy. US patent specification No. 5,119,409 proposes the use of dynamic pulse fluoroscopy with optimisation of beam quality and pulse rate but gives no description of the methods to be employed. US patent specification No. 5,845,269 proposes the use of a fuzzy logic control system for use in an X-ray diagnostics application.

Algorithms exist for independently controlling certain process parameters. For example, automatic exposure control tends to be based on brightness in either the whole image or in (manually) selected regions of the image. Generally, a photosensor is used to determine the brightness to input into a feedback loop that controls tube current. In X-ray imaging, because system magnification and image zoom are both adjusted by the operator, another control algorithm has been devised for automatically adjusting X-ray collimators to define the maximum field size so that the image area at the X-ray sensor is just covered.

The present invention has the objective of employing radiation-generated images from a monitored process in the direct control of the monitored process. It has the further objective of reducing the exposure to imaging radiation encountered by equipment or plant operators and, in the case of medical applications, to patients.

According to the invention there is provided a method for an industrial or medical process which employs radiation imaging to monitor progress of the process, in which a combination of parameters derived from the image is used to calculate output signals to enhance control of the process.

The invention offers the benefit that by feeding derived and calculated data back to the process a very high degree of control can instantaneously be achieved. The invention thus provides that the parameters of the imaging system can be controlled and altered dynamically based on information contained in and extracted from the images themselves.

In the medical field, for example in interventional neuroradiology using X-ray imaging, the invention offers the advantage that significant dose reduction is achievable compared with conventional X-ray imaging, thereby providing for optimisation of the minimum patient dose while maintaining diagnostic quality.

The parameters to be taken into account in the control calculations can include not only major "high level" parameters such as the intensity and duration of the applied radiation, movement of the target being monitored, composition and physical state of the target, but also secondary "low-level" parameters such as the absolute position of the target and its position relative to other elements. Other possible low level parameters include the absolute and relative velocity of the target, the uniformity of the target, and the image texture, pixel intensity and pixel noise.

The invention offers the advantage that calculations based on certain parameters monitored by the image permit the construction of other relevant parameters, for example the movement and composition of the target.

In one advantageous embodiment of the invention parameters derived from the image can be combined with other process data, for example ambient temperature and pressure observed by point sensors, to achieve further precision of control.

The data from the image and any related parameters can be converted into process control by appropriate image manipulation apparatus. In many instances it is desirable to employ suitable algorithms to form appropriate control outputs. Possible algorithm designs include rule based logic, fuzzy logic, neural networks and other linear or non-linear combinations. Algorithms can also be provided for cross-correlation of different parameters. A particular benefit is that cross checking parameters, even low level parameters, can ensure that the input data is robust. Similarly cross checking of outputs is beneficial in ensuring that that they are consistent.

Feedback of monitored data allows the use of algorithms to test predicted response against actual response and thereby permit control of oscillations

or unexpected results. It may however be desirable to delay a control adjustment to ensure that the monitored response is real and not simply caused by image noise.

Adaptive manipulation of input data (in time and space) may be of advantage to ensure the return of the most significant data for the particular control application.

Many other manipulations may also be beneficial within the scope of the invention. Images may be segmented to identify specific classes of input signal. Images may usefully provide functional data (for example in localised MRI spectroscopy, ultrasound Doppler signals and X-ray diffraction). Sensors can be dynamically adjusted in response to ambient conditions, even across the surface of the sensor (i.e. to accommodate wide dynamic range signals).

The type of image display is not critical. Images within the scope of the invention can incorporate stereo and 3D sound, colour overlay to indicate spatially localised parameter values, overlay of images with physical models of processes, overlay of multiple image data sets and production of statistical representations of data (e.g. 3D texture maps).

A method for addition of colour to greyscale images obtained from X-rays to facilitate inspection by an observer is described and claimed in our co-pending patent application of even date. A further co-pending patent application of even date relates to an improved collimator to control X-ray dose levels.

In a particularly beneficial embodiment of the invention the entire set of monitored parameters from the image and elsewhere can be analysed and manipulated to study an entire process with a view to achieving their efficiency improvements.

The invention is especially suitable for use in X-ray imaging using adaptive image processing. It is of particular benefit for medical X-ray fluoroscopic procedures such as interventional neuroradiology, cardiology and peripheral vascular angiography, and is hereinafter described mainly with reference to these procedures. It is however also applicable with advantage across a wide range of other medical procedures and industrial tasks.

X-ray fluoroscopy is a commonly used procedure for guiding interventional procedures within the body, or for visualising the structure/function of internal organs in the body. It is characterised by the use of X-ray imaging at video rate (normally 6 to 30 frames per second).

Conventionally, an X-ray imaging system for fluoroscopy comprises an X-ray irradiation unit (for example an X-ray tube and generator, collimator assembly, beam filter(s) and light beam diaphragm) combined with an imaging chain (for example, an X-ray image intensifier, lens system with optical iris, video camera, image processor and monitors). The images are observed by one or more specialist clinicians. Conceptually, the present invention combines a conventional X-ray imaging system with a new data processing apparatus (a "black box") which dynamically and automatically controls the operation of the X-ray imaging system based on the image data itself.

Most practical systems allow manual adjustment of some system parameters (e.g. frame rate, collimator position, display settings and tube voltage).

The improved system according to the invention comprises an "imaging computer" that interprets the video X-ray image and uses information derived from the images to drive such parameters as the position of the X-ray collimators, the X-ray pulse duration, the X-ray pulse frequency,

the X-ray tube voltage and the X-ray tube current. Additionally, the imaging computer can automatically adjust the displayed image by using colour, by stitching together live images with static or partially illuminated background images and by using temporal and/or spatial filtering. The imaging computer achieves this by extracting low level parameters from each image such as catheter tip position, local or global movement vectors, noise metrics and contrast/brightness data. The low level data are input to a predictive algorithm to ensure appropriate optimisation of the X-ray system for the next video frame.

With a view to ensuring the optimum selection of control algorithms the system can advantageously be programmed to allow an operator to enter fixed parameters of an individual field of activity, for example: the type of study (cardiac or neurological); whether the patient is on a moving couch; whether the main collimator is fixed or mobile; whether a contrast agent is being employed; whether the operator requires subtracted data, road map, live image etc.; the position of any reference/mask subtraction image; whether colour is to be used in the image and if so which colours.

Information drawn from each X-ray image in the fluoroscopy sequence is used to predict how the X-ray imaging system should be optimised for the next X-ray exposure. The result is thus image driven, giving dynamic optimisation and control of the entire X-ray imaging chain. The signals used to control the X-ray system are derived from adaptive (spatially and temporally variant) analysis of the X-ray images produced by the imaging chain. This approach maintains or enhances the clinical efficiency, namely the ability of the clinical team to perform their required task.

In one embodiment the invention thus provides that all the chosen system parameters may be controlled together, and not just a subset. A particular advance is the extraction of a set of low level parameters and their use together in a predictive optimisation algorithm. A particular benefit of

the invention in its application to fluoroscopy is its capability to minimise the X-ray dosage to reduce potential detriment to the patient and to the clinicians conducting the procedure.

A typical X-ray imaging system that uses control according to the invention in order to achieve radiation dose reduction combines a "black box" multi-processor data processing system with an X-ray generating apparatus and X-ray imaging apparatus. The X-ray generating apparatus may operate in a continuous output mode or in a pulsed fluoroscopy mode (i.e. using a pulsed X-ray beam with its intensity modulated as a series of discrete pulses). It includes X-ray attenuating diaphragms (collimators) that may be moved independently under computer control. In addition, the system should incorporate variable filters that can also be manipulated under computer control. The X-ray imaging chain should produce images with as little persistence and lag as possible so that the appropriate dynamic performance can be generated by the data processing system based on the information contained in the image sequence itself.

To achieve X-ray dose reduction in controlled fluoroscopic X-ray procedures, a variety of parameters may be varied. In the case of a "continuously on" beam of X-rays these are:

the area of the patient being irradiated,

the shape of the region of the patient being irradiated,

the tube voltage of the X-ray tube,

the tube current of the X-ray tube,

any filtration applied to the X-ray beam,

the aperture of an iris placed in the optical imaging chain (if using an X-ray image intensifier),

the gating time per frame of the image sensor,

the image lag introduced and controlled by the sensor electronics,

the gain and offset of the individual picture elements (pixels) of the image sensor,

the electronic gain and display contrast of the imaging chain (including the monitor),

the frame rate of the displayed X-ray image,

the ambient lighting conditions, and

the observer viewing distance.

In the case of a pulsed fluoroscopy system additional parameters for control are:

the pulse rate at which individual X-ray exposures are delivered to the patient,

the width of the individual X-ray pulses,

the X-ray tube current and voltage profile during the X-ray pulse, and the relative rates of X-ray pulsing and display updates.

By dynamically altering some or all of these parameters, a significant reduction in dose may be obtained without affecting clinical efficiency.

In clinical X-ray imaging, examination of the statistical values of the parameters collected over an entire sequence (e.g. area of patient being exposed, X-ray tube pulse rate, kV, mA etc) can provide a comparison with other sequences, for example in assessing individual operator performance. Further, the sum of collected data can determine where the

major dose delivery is occurring, thereby enabling improved dosing to be planned and implemented.

Low level parameters may be extracted from the fluoroscopy image sequences for use in real-time optimisation of the X-ray imaging system, with the particular aim of reducing patient dose. Low-level parameters in fluoroscopy include the position of a catheter tip, the position of a catheter guide wire, the location of vessels containing contrast agent and the position of the X-ray collimators. Further relevant information may include local movement vectors for extracting anatomical motion (e.g. rate of flow of contrast agent in a particular vessel), global motion vectors (e.g. change in patient position with respect to the X-ray source) and image statistical properties (e.g. brightness, signal-to-noise ratio, noise-power spectrum and image contrast). To achieve dynamic optimisation, frame-by-frame evaluation of these parameters is required. Generally, such methods will have relatively low accuracy and/or reliability. Therefore, in order to ensure efficient feature extraction, the frame-by-frame methods are mirrored by other, potentially more accurate and sophisticated but slower, methods to check/re-align the frame-by-frame methods. This ensures overall system reliability.

The invention thus provides for low level parameters to be extracted from X-ray images for the purposes of dynamic optimisation of the X-ray system. Further, these parameters may result from quick but non-robust methods operating within a single video frame or from sophisticated methods that update less frequently. Parallel computation may be employed to achieve sufficient speed for image analysis.

Low level parameters extracted from medical fluoroscopy image sequences may be combined by high-level algorithms for real-time optimisation of X-ray imaging systems with the aim of reducing patient dose. High level algorithms include linear rule based logic, neural

network based methods, fuzzy logic and statistically based algorithms. Two outputs are required from these algorithms: hardware optimisation and image display optimisation. The hardware algorithm should control tube voltage, tube current, X-ray pulse frequency, X-ray pulse duration and collimator aperture. The image display optimisation should control spatio-temporal filtering, the use of colour, histogram equalisation and unsharp masking methods.

Thus the invention further provides for low level inputs to be fed to control algorithms to undertake optimisation for system control and optimisation for image display. Further, the invention permits partitioning of these algorithms into parallel computing systems.

The data processing apparatus comprises a computing system programmed with algorithms required to implement the appropriate control and dose reduction strategies. Each algorithm is best suited to a particular type of procedure (e.g. cardiology, barium fluoroscopy etc) and the data processing apparatus therefore needs to contain multiple algorithms, including one for each type of procedure. Each algorithm is based on observation of the approach taken to the clinical task by the human observers performing the procedure, since it is the human perception of the quality of the image sequence that is the main criterion used for minimisation of the delivered dose.

A patient undergoing treatment is typically on a movable couch beneath a stationary X-ray irradiation unit. Movement of the couch effects global motion of the patient with respect to the irradiation unit. During this period of motion, the human visual system is unable to perceive fine spatial details because they are moving too fast. In order to maintain reasonable statistical accuracy in the displayed pixel values it is therefore preferable to reduce beam current and consequently coarsen the pixellation of the displayed image. Further, because the processing

apparatus can detect the magnitude of movement from one frame to the next, it can use this information to re-register into the displayed image one or more "good" frames taken when the patient was static before the couch movement. This can minimise the perceived effects of image lag on the image intensifier. Additionally, the processing apparatus can detect the acceleration of the couch, and hence determine how long the movement is likely to continue. This information may be used to modulate the pulse rate of the X-ray beam, hence allowing frames to be dropped where couch movement is greatest and the quality of the recorded image is naturally lowest (due to image lag and image smearing during a given X-ray exposure). Compared with a conventional system the combination of these three procedures allows a reduction in dose of at least an order of magnitude during the couch movement.

Alternatively, during interventional procedures using a catheter, the attention of the clinical operator is generally focussed on the site at the tip of the catheter. Therefore, by using dynamic collimation of the X-ray beam, it is possible to track the region around the tip of the catheter with high X-ray exposures (hence good feature visibility), while delivering zero or low dose to peripheral regions of the image. This requires the displayed image to be constructed from the "live" catheter image plus a background image constructed from image data acquired with a previous full frame exposure or from multiple frames acquired at low-dose that are re-registered to fit around the "live" catheter image. Additionally, it is possible to modulate the frame rate of the high dose localised catheter image based on the rate of movement of the catheter within the image. Combining these effects leads to high levels of dose reduction in many interventional situations.

To minimise distraction to the clinical operator, it may be desirable to add computer generated noise into the background image. Surprisingly the added noise can add to the visibility of features in the background

region. This occurs since the detection of low contrast features by the human visual system can be enhanced by small temporal changes in image brightness within the static feature.

In many embodiments of the invention it is advantageous to use graded image filtration by inserting shaped filters into the X-ray beam. Suitable shapes of filter include linear wedge, exponential or parabolic profiles. With a wedge-shaped filter X-ray exposures are greatest at the (thin) tip of the wedge, and lowest in the thick regions of the wedge. By using two or more wedge-shaped filters, it is possible to reduce the dose significantly in the peripheral regions of the image, while maintaining full dose levels in the critical diagnostic regions of the image. Movement of the filters can be controlled by a suitable algorithm.

Adaptive image processing techniques are preferably used to perform temporal averaging with re-registered image frames in low-dose background regions while unmodified image data is displayed in high dose regions. When combined with image driven dynamic movement of the filters, this approach can lead to major dose savings

An electronics system may be used to apply algorithms for real-time optimisation of X-ray fluoroscopy systems for minimisation of patient dose. Such systems generally require parallel electronics designs and are required to provide means for image acquisition, image display, low level parameter extraction, high level algorithm implementation and external X-ray system control. The system according to the invention may be constructed using a number of technologies including multi-processor commodity component systems (based on PC/DSP/PGA approaches), beowolf class supercomputers or SIMD/MIMD supercomputers, custom gate arrays, custom integrated circuits, custom processors and/or three-dimensional interconnect solutions. Algorithms are segmented to run on parallel systems e.g. calculation of each low level parameter is assigned

to one or more processors, these algorithms being run in parallel to provide data for a concurrent high level algorithm, which in turn controls a parallel image display algorithm. Use of pipelining and highly segmented memories is usually required in such systems.

According to a further useful embodiment the invention provides a method for construction of an image computer specifically for dose reduction in fluoroscopy.

The algorithm required for each class of clinical procedure will normally be different. For example, an algorithm designed for peripheral vascular procedures will not normally be successful for cardiac imaging. Further, there is additional information that is normally generated by the operator in terms of the type of display they would like to see. For example. the clinician may wish to see a live projection image or a subtraction image for visualisation of contrast agents.

The operator also provides input to the imaging system, for example in terms of mechanical movement of the patient couch or other equipment with respect to the patient. These user generated input parameters are used by the optimisation algorithms alongside the low and high level parameters generated directly from the image.

To ensure efficient algorithm implementation, it is normal to provide an operator interface to the dose control system. This normally contains one or more buttons and one or more indicators and may include a text or graphic display.

One of the buttons normally fulfils the function of system override. On pressing this button, the X-ray system returns to a normal, typically preset mode of operation. When the clinician releases the button, the automatic dose control system will either take over directly or pause for a preset time before resuming operation.

A further button or set of buttons may be used for technique selection, for example to select between cardiac or peripheral vascular studies. Alternatively. this information may be provided by a data link to the main X-ray system.

A further button or set of buttons may be used for algorithm control. For example target dose reduction levels may be requested as being "high or "normal".

Further buttons may be used to control other aspects of system operation as required.

Typically, the dose control apparatus will include an electronic link to the X-ray diagnostic system through which the output control signals are propagated. This electronic link will normally also contain inputs related to parameters such as couch movement or movement of the X-ray system relative to the patient. It may also contain input signals to indicate the injection of contrast agent.

Indicators may be provided to confirm the status of user specified input parameters, for example the class of the requested clinical technique.

As a further example, a text or graphic display can be provided to indicate the actual setting of the control outputs. For example the display may indicate, for example in either text or graphical form, the tube voltage. tube current, pulse rate and cumulative procedure dose.

By collating statistics obtained during the procedure. it is possible to compare the administration of a particular treatment with other similar treatments. Further, it is possible to compare the statistical properties of a group of treatments. Such information may be used to plan enhanced treatment protocols, for example to minimise patient dose and/or minimise overall procedure time. Such information can both be used to plan optimal treatment protocols and optimal utilisation of equipment and staff.

Suitable statistical properties that may be generated automatically include, but are not restricted to, the value of all high and low level input parameters and the value of all control outputs. For example, it is possible to collate information regarding X-ray field size, instantaneous dose rate. time into the procedure. X-ray tube pulse rate. tube voltage and current, velocity of catheter, position of catheter tip with respect to landmarks and so on.

The database formed by this information may be structured such that data may be retrieved as a function of any one or more of these values. Results from a search of the database may be structured in many ways including textual, graphical, tabular. numerical or image format. Colour may also be used to highlight specific data.

For example, a control image or scatter plot may be generated from clinical image data to show the correlation between two or more variables, for example between dose rate and catheter tip velocity. Colour can be applied to this image to indicate regions where, as an example, tube voltage increased above a certain limit, say 100 kVp. As a further example, colour can be applied in the image to indicate uncertainty in catheter tip position. A typical use of this information is to support algorithm testing.

As a further example, a graph may be constructed of dose rate versus procedure time. Colour may bc used to indicate for example catheter tip velocity, injection of contrast agent and/or X-ray tube pulse rate all on the same graph. It is then possible to analyse the data to observe such parameters as the fraction of the procedure for which high doses are being used, whether the catheter is moving substantially during this time and how frequently contrast agents are injected. This information is useful to the clinician in terms of review of their clinical practice.

By performing statistical tests of data from one procedure against data from other similar procedures, it is possible to build up a generalised model of particular classes of clinical procedures.

As an example of the use of such a model, it is then possible to assess individual clinician performance against the generalised model procedures. Such results may be used to aid initial training and for continuing education of the individual clinician.

As a further example, models representing the collective performance of individual institutions may be compared against national models derived from data produced by sets of institutions. Such comparisons can result in improvements in clinical practice and national standards for clinical protocols.

As a further example, national models may be compared against other national models to develop internationally acceptable clinical protocols.

The invention is further described with reference to the accompanying figures, in which:

Figure 1 is a schematic view of an imaging system according to the invention;

Figure 2 is a diagram showing in simplified format the calculations which need to be performed within the system of Figure 1;

Figure 3 is a diagram showing in simplified format the structure of the data processing required in the system;

Figure 4 is a diagram showing in simplified format a commodity computer cluster for the required data processing;

Figure 5 is a diagram showing in simplified format a multi-processor bespoke cluster for the required data processing;

Figure 6 is a schematic view of the data processing portion of the system.

In the schematic view of the system given in Figure 1, an X-ray imaging chain feeds received signal data to an image processing apparatus which in turn controls an X-ray irradiation unit. The image processing apparatus however also feeds back processed data to the X-ray imaging chain such that the latter supplies to one or more clinical observer(s) not only information directly from the image but also the said information as enhanced by algorithms. The clinical observers have control over the image processing apparatus but this control is effected with the enhanced information from the imaging chain.

The calculations which need to be performed within the image processing apparatus are summarised by the simplified diagram given in figure 2. This shows that processing needed to optimise the display of a new image frame is not necessarily the same as the processing required to predict the optimal state of the imaging system for the next image frame.

Typically there is a single video frame delay between the acquisition of the new image frame and its display. For a video standard running at 25 frames/second, this constitutes a delay of 40 ms. However since this is below the threshold for temporal response of the human visual system it does not represent a significant detriment to clinical judgement.

Figure 3 summarises the structure of the required data processing. Regardless of the algorithm being used, the underlying processing required is similar. In principle, a given algorithm is relatively straightforward to compute with a set of control outputs (diaphragm position, tube voltage, frame rate etc) and a set of optimisation inputs derived from the X-ray image data (catheter position, movement data, noise information etc). The algorithm itself may use one or more computational techniques such as rule based logic, fuzzy logic, neural networks or other linear on non-linear combination methods to combine

these inputs to generate the outputs. The calculations can be implemented extremely rapidly using modern computing methods.

In the case of a rule based logic system, inputs are combined in standard ways and compared against reference values determined by clinical trials. Generally, multiple "low-level" inputs are combined into one "high-level" input, the values of the high level inputs being weighted to determine the appropriate control output. The values of the weighting factors are normally determined through clinical trials.

For example, suppose that A, B and C are low level inputs related to movement while D and E are low level inputs related to noise. Suppose that high level input X is the result of combining A, B and C while high level input Y is the result of combining D and E. The impact of these high level inputs on control output 1 (e.g. O1, X-ray tube pulse rate) is determined in the form $O1 = x_1X + y_1Y$ where x_1 and y_1 are reference values for controlling O1. Different combinations of reference values and high level inputs will generally be used to determine the optimisation strategy for each of the other control outputs.

A second level check is preferably performed to ensure that conflicts between output control settings are eliminated. For example, if the noise high level input suggests that more X-ray photons are required to form an acceptable image, this can be achieved by increasing tube pulse rate (O1), increasing tube voltage (O2), increasing tube current (O3) or undertaking more extensive temporal averaging in the displayed image (O4). The appropriate strategy is determined by further linear combination of factors (e.g. select greatest of [aO1, bO2, cO3, dO4] where a, b, c and d are constant multipliers). In this example, increasing tube current is likely to be the preferred option over changing tube pulse rate, changing tube voltage or changing temporal averaging (and hence $c > (a, b, d)$).

In a preferred extension of this approach, an iterative procedure is undertaken in which new values of the control outputs are selected and fed back into the algorithm to test the likely outcome in the next image frame. The process is repeated until the values converge to a suitable limit.

Generally, the output control values are constrained to prevent unacceptable fluctuations in system performance (e.g. rapid variations in image brightness or signal-to-noise ratio).

A further algorithm approach uses statistical testing to evaluate the values of low level and high level inputs and of each control output and the set of control outputs. For example, if a high or low level input value moves, or is predicted to move, significantly outside an expected range, the algorithm should determine, and test, a suitable strategy for bringing the value back into the acceptable range. It does this by ensuring that the individual output control values all remain within their expected ranges, and also that the set of control outputs remain within their expected range. For example, kV and mA have individual ranges while their product, mA x kV = kW, has a separate allocated range.

The expected range of all values is determined through clinical testing. Generally, an iterative approach is adopted whereby the output control values are adjusted in order to bring the required input value back into range without driving a different input value out of range.

An optimisation algorithm may also be developed based on the neural network principle, whereby the network is trained by classifying the results of the algorithm as being clinically "acceptable" or "unacceptable". The algorithm requires input data that is derived from both the most recently acquired X-ray image, and also from the set of previous X-ray images.

It is generally much more computationally demanding to acquire the information required by the control algorithm from the input image data than it is to implement the algorithm itself.

It is convenient to visualise the input data as being of "high value", having been assembled from several "lower value" parameters. For example, as indicated in figure 3, the high value input "movement" can be constructed from several lower value data parameters including for example a map of local image movement vectors and a single global image movement vector. In this way, the low level data used by a given algorithm has previously been subject to strong filtering specific to the design of the algorithm itself.

In some situations, low lèvel data is fed to the algorithm directly. In interventional procedures, the location and direction of motion of the tip of the catheter are highly significant and together form a key component of the input data required to make a decision.

Generally, it is best to generate both the high level and low level input data in a number of ways, so that one method can "cross-check" the output produced by the other. For example, a quick but potentially non-robust method for locating the catheter tip may fail if the catheter makes an unexpected movement, while a more sophisticated, but time consuming method will track the catheter at all times. Therefore, the accurate, but slow, method is used to correct and/or cross check the rapid but less accurate method on a periodic basis. This may be extended to multiple methods, each correcting/cross checking each other at suitable time periods. A statistical check on the output of each method is used to determine confidence in the overall result for the parameter. If confidence falls below an acceptable value, optimisation of the X-ray system is normally automatically terminated and the entire X-ray system returned to standard dose operation as soon as possible. Once the results of the

methods return into the confidence region, optimisation of the system can restart.

An example of such an approach is in catheter tracking. A quick but non-robust approach selects a small region of interest, say 30 x 30 pixels, around the catheter tip. When a new image is generated, the data in the region of interest is subtracted from the data in the region of interest of the previous frame. The subtracted data is compared on a pixel by pixel basis, and any pixels with a statistically significant value from the mean value (e.g. 2σ from the mean level) are assumed to be due to movement of the catheter. Therefore, using the sign as well as the magnitude of the subtraction data, the new catheter position can be rapidly updated.

Based on the movement detected, it may be necessary to modify the shape and/or size of the region of interest of a frame-by-frame basis to ensure reasonable catheter tip detection. However, for this method to work, it is necessary for the region of interest to be placed around the catheter tip in the first place. This requires a more sophisticated algorithm, for example a segmentation-correlation method, Bayesian estimator approach or other suitable method. Hence the sophisticated (but relatively slow) method is used to support the fast, but uncertain accuracy, method.

If the catheter tip position generated by the two methods is found to be inconsistent, then the confidence in this parameter is assumed to be reduced. This would occur, for example, if significant patient movement occurred between two image frames. The control algorithm then evaluates whether to return to normal (i.e. non-optimised) operation for a short period of time until confidence in the parameter values is restored.

Low-level parameters that are usually important in medical X-ray fluoroscopy include a binary segmented map of vascular structure, the velocity of contrast flow in blood vessels, the structure and texture map

of vessel diameters, local image statistics (to help locate image features), local image movement vectors, global motion vectors, and the position of anatomical landmarks in the image (for help with image registration),

As well as optimising the control of the imaging system, it is also necessary to generate a clear display for the image. Here, the high level information used by the control algorithm may also be used to help image display.

For example, the vascular structure determined as a low level data set during a previous injection of contrast agent can be displayed over the live image after re-registration for any patient movement. This can reduce the quantity of contrast agent injected into the patient while also minimising procedure time since the clinician can see, not simply remember, where the vascular structure is. Generally it is preferred to display the vascular structure as a series of line segments rather than true shapes to ease visibility of the underlying anatomical structures.

A further example is the use of adaptive temporal and spatial image processing. In particular, temporal averaging on a pixel-by-pixel basis depending on the low-level movement information for that particular pixel; where movement is significant, do not apply temporal averaging, but where movement is insignificant, apply temporal averaging. Similar methods can be used for spatial smoothing; apply smoothing over regions of insignificant contrast, but do not apply smoothing to regions of significant contrast.

A further example is the use of frame summing, where images collected previously may be re-registered and summed with the current image to reduce effective photon noise, while maintaining image sharpness.

A further example is where reference images collected in the past may be used to back-fill new images where part of the new image is obscured by a collimator.

A further example is where the background image can be updated by regions of the new image which are not obscured by the collimator (i.e. those parts of the most recent image that contain good data).

A further example is where the partially attenuated live image is summed with the background image to boost the contrast in the composite image compared to the true live image (i.e. if background image = A and live image = B, then the displayed image = aA + bB where a and b are constants).

A further example is where adaptive contrast stretching is used to maximise visibility of features in the displayed image. A particular method, for example, is histogram equalisation.

A further example is where unsharp masking, e.g. superimposing a blurred image on a sharp image, is used to improve visibility of features in the displayed image.

A further example is where the true pixel values are scaled to brighten or darken a displayed image to achieve normal brightness. This is particularly useful in correcting the displayed brightness of pixels obscured by the partially attenuating collimator.

A further example is where random noise is added to the background/reference image in order to improve perceived quality and/or uniformity of the displayed composite background/live image.

A further example is where the displayed pixel dimension is altered depending on the value of localised low level parameters such as movement at the pixel level. In regions of low movement, pixel sizes are

increased (e.g. 2 x 2 or 3 x 3 pixel regions) by taking the average or median value to enhance signal-to-noise ratio, while pixels are displayed at the normal size in regions of significant movement or high contrast. The subtly alters the user perception of the displayed image and can substantially improve observer performance.

The data processing apparatus described above should be capable of sustaining high data rates with excellent compute capability. In general, this requires the use of parallel computing systems. In principle this is most likely to be achieved in one of the following ways:

(1) by using a single highly specific processor with multiple internal processing elements and data busses to perform all processing,

(2) by using a cluster of commodity computers connected through a local area network or high speed backplane, or

(3) using a bespoke multi-processor compute system.

It is not normally feasible to design a single processor to perform the algorithms to work at video frame rates. An example configuration for a commodity computer cluster is shown in figure 4 while an example multi-processor bespoke system is described in figure 5. Figures 4 and 5 are representative of other suitable configurations, although many other configurations are possible. In particular, three-dimensional interconnects, programmable logic devices and multi-chip modules can all be used effectively for implementing parallel algorithms.

In multi-processor configurations, there are many ways to share the computational load. When using a cluster of computers, network bandwidth tends to dominate the overall system throughput and it is then efficient to segment the primary image and to send one small segment to each node in the cluster. Each node then processes the image segment to

extract all relevant low level information. When this data has been obtained, the data is transferred to a different node to be assembled with data calculated for the other image segments to form high level information. The high level information is then transferred to a further compute node to implement the system optimisation algorithm. It is also efficient for the cluster to use multi-processor nodes. In this case one (or more) processor(s) at each node can generate the low level data required by the optimisation algorithm while one (or more) other processor(s) can generate the new output image data. Here, the output image data is calculated on the segment of image data used previously to generate the low-level control information and hence network traffic is minimised.

A multi-processor bespoke system will typically copy the full image to a number of processors simultaneously through a high bandwidth backplane. Each processor then extracts a given high level parameter (e.g. movement, noise) from that image. The high-level data is then transferred to a master processor for implementing the system control algorithm. Alternatively, image data may be segmented and processed as described previously. Typically, each processor is based on a high-end digital signal processor (DSP) that may itself contain eight or more processing units and multiple internal and external memory busses. Where appropriate, DSP and memory integrated circuits are combined with high performance programmable logic circuits to implement what might otherwise be computationally intensive tasks (e.g. binary segmentation and/or convolution). In this way, real-time performance can be achieved within a reasonably sized package.

The dose reduction strategy described here may be implemented on standard or existing equipment by retrofitting suitable beam diaphragms and variable filters and adding an image data processing system. A schematic of a possible data processing system is given in figure 6. Ideally, the existing system would use a pulsed X-ray tube, but significant

dose reduction can still be achieved using a standard continuous output X-ray tube.